

ABSTRACT

A method and apparatus are disclosed for allocating functional units in a multithreaded very large instruction word (VLIW) processor. The present invention combines the techniques of conventional VLIW architectures and conventional multithreaded architectures to reduce execution time within an individual program, as well as across a workload. The present invention utilizes a compiler to detect parallelism. The disclosed multithreaded VLIW architecture exploits program parallelism by issuing multiple instructions, in a similar manner to single threaded VLIW processors, from a single program sequencer, and also supports multiple program sequencers, as in simultaneous multithreading. Instructions are allocated to functional units to issue multiple VLIW instructions to multiple functional units in the same cycle. The allocation mechanism of the present invention occupies a pipeline stage just before arguments are dispatched to functional units. The allocate stage determines how to group the instructions together to maximize efficiency, by selecting appropriate instructions and assigning the instructions to the FUs. The criteria for selection are thread priority or resource availability or both. Under the thread priority criteria, different threads can have different priorities. The allocate stage selects and forwards the packets (or instructions from packets) for execution belonging to the thread with the highest priority according to the priority policy implemented. Under the resource availability criteria, a packet (having up to K instructions) can be allocated only if the resources (functional units) required by the packet are available for the next cycle. Functional units report their availability to the allocate stage.